

Aid, Growth, and Poverty

Aart Kraay
The World Bank¹

*Presentation for IMF Seminar on Foreign Aid and Macroeconomic Management
Maputo, Mozambique, March 14-15 2005*

1. Introduction

The organizers of this conference have asked me to speak today about aid and poverty alleviation. We have already heard Steven Radelet provide a very thorough review of the state of the evidence on the effects of aid on growth.² I would like to use this presentation to make three broad points linking this evidence on aid and growth to poverty alleviation.

First, I would like to review the cross-country evidence on the links from growth to poverty reduction, in order to underscore a basic point that should not be too controversial: sustained poverty reduction is impossible without sustained growth. In fact, I want to go a step beyond this and argue, based on cross-country empirical evidence, that the recent vogue among development practitioners for debating the extent to which growth is “pro-poor” may be somewhat misguided. As we shall see cross-country differences in the extent to which growth is more or less pro-poor are dwarfed by cross-country differences in growth performance itself. This should focus our minds squarely on understanding the fundamentals of growth as a vehicle for poverty reduction.

Second, while I believe that there is plenty of evidence in support of the proposition that aid has a positive impact on growth (and through this, on poverty), I want to put the magnitude of the growth benefits of aid in perspective. In particular, while several careful and sensible papers have been able to find positive and significant effects of aid on growth, it is important to note that the share of the cross-country

¹ 1818 H Street NW, Washington, DC 20433, akraay@worldbank.org. The views expressed here are the author's and do not reflect the official views of the World Bank, its Executive Directors, or the countries they represent. I would like to thank, without implication, Alan Gelb for helpful comments.

² Clemens, Radelet, and Bhavnani (2004).

variation in growth performance that we can explain using aid is typically not that large. Measures of policies and institutions, as well as exogenous shocks, account for much more of the cross-country variance in growth than does aid itself. And of course, as is common in the entire empirical growth literature, there is a very large share of growth that remains unexplained.

This brings me to my third broad point: gaining a better understanding of the non-aid determinants of growth, and the extent that foreign aid can help to improve them, seems important for thinking about the effectiveness of aid in the long term as an instrument for poverty reduction. This point is of course very broad and potentially opens the door to a discussion of very many issues. Since there is no way I can be comprehensive, I will instead be very selective and focus on a two issues that I think are important and that I have been working on in my own research.

One of these is the link between good governance and economic performance. We by now have a great deal of empirical evidence showing an important causal effect of better governance on economic outcomes. Against this background, the on average poor performance of countries in Sub-Saharan Africa on most measures of governance, and the absence of any clear evidence of improvements in recent years, is of great concern. Yet some of the recent discussion of scaling up of aid to Africa seeks to minimize governance issues in the region by pointing out that governance is no worse than would be expected given low income in the region. Claims like this are difficult to substantiate without a clear understanding of the directions of causation between income and governance, and vice versa. As I will argue in more detail below, I do not think that there is much compelling evidence to suggest that weak governance in Sub-Saharan Africa can be attributed to low income, and so it think it is misleading to suggest that governance is not an issue in the region. And a better understanding how aid should be allocated to recognize differences in governance that affect how well it can be used, as well as a better understanding how aid can be used to improve governance, seems central to the growth and poverty reduction agenda in the region.

The other issues concerns the importance of poverty traps for understanding low income levels and the slow pace of development in Sub-Saharan Africa. Many observers have argued that a variety of self-reinforcing mechanisms, or poverty traps,

are responsible for the region's stagnant income levels. To the extent that this argument is correct, it has important implications for aid and for development policy. Put simply, the existence of poverty traps lends support to the idea that a "big push" or a major scaling up of aid is required to break countries out of these poverty traps and launch them on sustainable growth paths. Given the significance of such a policy recommendation, it seems important to have a good empirical understanding of the importance of such poverty traps. But as I will discuss in more detail below, in my view there is actually relatively little compelling evidence that such traps exist. This is not to say that substantial increases in aid to Africa are a bad idea -- to the contrary there probably are very many useful interventions and policies that foreign aid can support. But the absence of clear evidence supporting poverty traps should caution us against thinking that large increase in aid are a necessary condition for growth, nor that large increases in aid will have proportionately much larger growth benefits than small increases in aid.

2. Growth and Poverty Reduction

Figure 1 shows the relationship between growth and poverty reduction, using household survey data from a sample of developing countries in the 1980s and 1990s. Each point in the graph corresponds to an episode or "spell" lasting at least five years and averaging 10 years in length, over which available data allows us to calculate the proportional change in the headcount measure of poverty and the proportional change in mean income or consumption. It is clear from this graph that there is a very strong negative relationship between growth and changes in poverty. Virtually all observations are clustered in the bottom-right quadrant (positive growth and declines in poverty) or the top-left quadrant (negative growth and increases in poverty).

Certainly there are also deviations from this average relationship. In Uganda, Mauritius, and Ghana, for example, growth was similar in the 1 to 3 percent range over the indicated periods, but the rates of change in poverty ranged from about -8 percent to +2 percent per year. For Africa as a whole, it is telling that 8 of 12 countries fall above the regression line, indicating a poverty reduction performance that was worse than what would be expected for a typical developing country with similar growth performance.

How should we think about deviations from this average relationship between growth and poverty reduction? To answer this, it is useful to decompose changes in poverty somewhat mechanically into three components:

$$\text{Percent Change in Poverty} = \text{Growth in Mean} \times \text{Sensitivity of Poverty to Growth} \\ + \text{Changes In Relative Incomes}$$

The first two terms in this decomposition, growth times the sensitivity of poverty to growth, are referred to as the “growth component” of changes in poverty. This captures the extent to which poverty would have fallen had there been no changes in relative incomes, i.e. no change in the distribution of income, over the period. The third term captures the part of the change in poverty that is attributable to changes in relative incomes.³

In a recent paper I have empirically implemented this decomposition using a large cross-country dataset on growth and changes in poverty (Kraay (2005)). A useful way to summarize the relative importance of these components of poverty reduction is by using variance decompositions. Figure 2 provides a graphical depiction of this variance decomposition. It graphs the total change in poverty on the horizontal axis, and the growth component of changes in poverty on the vertical axis. The slope of the line of best fit has a variance decomposition interpretation: it is the share of the cross-country variation in changes in poverty that can be accounted for by cross-country differences in the growth component of changes in poverty. The fact that this slope is 0.97 indicates that virtually all of the cross-country variation in changes in poverty is due to cross-country differences in the growth component. Conversely, virtually none of the variation in changes in poverty is due to changes in relative incomes. Put simply, the graph shows us that if poverty fell, it is by far most likely that this is because the growth component of changes in poverty was large (in absolute value), and it is quite unlikely that poverty fell because income inequality fell in such a way as to reduce poverty.

We can also decompose the growth component of changes in poverty into growth itself, and cross-country differences in the sensitivity of poverty to growth. While it is more difficult to do a variance decomposition here (because we now have the

³ This decomposition was introduced by Datt and Ravallion (1992).

product of two terms rather than the sum), it still is very informative to graph the growth component against growth itself, as is done in Figure 3. It is clear from this graph that if the growth component of poverty reduction is large, it is most likely that growth itself was large, rather than that the sensitivity of poverty to growth was large. This graphs captures the fact that cross-country differences in the sensitivity of poverty to growth are in fact relatively small.

The variance decompositions in the two previous graphs refer only to the headcount measure of poverty, and also to a set of spells of poverty changes that cover fairly long periods of time, averaging around 10 years. I have also implemented these decompositions for different poverty measures, and for a much larger sample of short spells of changes in poverty averaging around three years in length. These results are summarized in **Table 1**. Briefly, I find that for more bottom-sensitive poverty measures, the share of poverty reduction due to growth is somewhat smaller than for the headcount. This evidence does not mean that the poorest of the poor are less likely to share in the benefits of growth. Rather it simply reflects the fact that more bottom-sensitive poverty measures place lower weight on growth in average incomes. It is also the case that over shorter periods the share of growth is somewhat smaller than for the sample of long spells. This to a large extent reflects greater volatility of measured inequality within countries, which tends to average out considerably over time. Overall however, these results are broadly consistent with the results we have seen for the headcount measure of poverty: most of the variance in changes in poverty is due to growth itself, and especially so over the medium to long run.

This decomposition is also useful for thinking about why it is that countries in Africa have on average had slower poverty reduction for a given rate of growth, as we saw in Figure 1. One possibility is that countries in the region were relatively more likely to see increases in inequality which dampened the effect of growth on poverty. We have already seen that on average, these relative income changes account for very little of changes in poverty. The top panel of Figure 4 plots this (typically small) component of poverty reduction on the vertical axis, against growth itself on the horizontal axis. It is interesting to observe that the spells in countries in Africa are fairly evenly distributed above and below the horizontal heavy line corresponding to the median for all countries.

What this tells us is that countries in this region were not especially likely to have increases in inequality offsetting the effects of growth on poverty reduction.

The other possible explanation for slower-than-expected poverty reduction given growth is that the sensitivity of poverty to growth has been lower in Africa. This is highlighted in the bottom panel of Figure 4, which plots the sensitivity of poverty to growth (holding constant relative incomes) on the vertical axis, and again has growth on the horizontal axis. Here the striking observation is that ten out of the twelve spells occurring in Africa fall above the heavy line indicating the median value for all countries. This means that the sensitivity of poverty to growth in Africa was on average lower (in absolute value) than in the rest of the world. This lower-than-average sensitivity of poverty to growth in Africa can be traced back to Africa's low income levels, and somewhat higher than average inequality, which together dampen the effect of growth on poverty. In my view this lower-than-average sensitivity of poverty to growth in the region also underscores the importance of more rapid and sustained growth for poverty reduction.

Finally, this empirical evidence is useful for the broader discussion around the notion of "pro-poor growth" which has become widespread in recent years. While the term "pro-poor growth" has become quite popular, there is as yet not much consensus as to what exactly it means. At one extreme is the view that growth is pro-poor only if relative incomes change in such a way as to reduce poverty (Kakwani (2000)). At the other extreme is the view that growth is pro-poor as long as poverty falls (Ravallion and Chen (2003)). It seems to me that the choice between these two definitions is fairly clear, and that the latter definition makes much more terminological sense. Consider the case of China's rapid growth, and also rising inequality, since 1980. This period has seen enormous reductions in the numbers of the absolutely poor in China, and it would seem quite odd not to refer to this experience of massive poverty reduction as being "pro-poor" simply because inequality also increased over this time. Conversely there have been cases of countries with declining incomes, rising poverty, but also declining inequality – it would also seem peculiar to refer to such episodes as "pro-poor".

There is however in my view a more substantive problem with this discussion of pro-poor growth. It seems to me that the value-added of the "pro-poor" prefix appended

to the term growth depends greatly on (a) the extent to which there are differences between pro-poor growth, and just plain growth, and (b) the extent to which policies can influence these differences. The variance decompositions that we have seen suggest that the distinctions between growth, and pro-poor growth, are actually quite small on average in the historical data. If we adopt the purely relative definition of pro-poor growth, then discussions of the “pro-poorness” of growth concern just a tiny part of poverty reduction. And even if we take the second and more sensible definition of pro-poor growth, the fact that the growth component of changes in poverty is dominated by growth itself (particularly in the long run) again suggests that the gap between growth, and pro-poor growth, is relatively small. Finally, at the cross-country level of analysis, we do not yet have a great deal of empirical evidence as to the policy determinants of changes in relative incomes that might make growth more pro-poor. There are relatively few robust cross-country findings on the determinants of levels or changes in summary statistics of inequality.⁴

In summary, we have seen that reductions in poverty require sustained growth, and that historically most of the variation in countries’ experience with poverty reduction can be traced back to cross-country differences in growth performance. This points to the centrality of growth for poverty reduction, which in turn underscores the importance of aid, and other factors, for growth. Before moving on to these issues, two final qualifications or caveats are in order. First, in all of this discussion I have focused on income and/or consumption-based measures of poverty, which as we have seen move very strongly with changes in average incomes. However, we do also care about non-income dimensions of poverty, notably health and education outcomes. While these are also strongly correlated income levels across countries, within countries over time the relationship between growth and improvements in health and education outcomes tends to be weaker.⁵ There are many reasons for this, and their discussion would occupy an entire other presentation. For now however I want simply to acknowledge that the links between growth and these non-income dimensions of poverty reduction are not as strong as the ones we have seen with income poverty. Second, the fact that on average

⁴ See Dollar and Kraay (2002) for an example of the difficulty in finding statistically significant determinants of the first quintile share using cross-country data. In addition, in Kraay (2005) I also look at the precise measures of inequality change that matter for different poverty matters of interest, and I again find no strong patterns linking these to a variety of measures of policy.

⁵ See for example Easterly (1999).

changes in relative incomes matter little for poverty reduction, and the fact that we have few robust cross-country correlates of these relative income changes does not mean that policymakers in any country can be oblivious to the distributional consequences of macroeconomic policies. Any policy change in any given country will have individuals differently, and careful country-specific analysis can shed light on likely impacts. But neither of these qualifications undermine the broader message of this section concerning the importance of growth for poverty reduction.

3. Growth and Poverty Reduction Depend on Much More Than Aid

In the previous section we have seen that growth is central to poverty reduction. We also have a large body of evidence pointing to a significant impact of aid on growth. As has been nicely summarized by Steve Radelet, this impact of aid on growth may depend on the type of aid, how it is financed, the time horizon, and also on the policy and institutional environment of the recipient country. Putting these two observations together provides a direct link from aid to poverty reduction. But at the same time I think it is important to put this channel from aid to poverty reduction in perspective, by noting the obvious: growth depends on much more than just aid.

This point can be seen most vividly by taking a representative growth regression from the aid and growth literature. I take one of the core specifications from Burnside and Dollar (2000), which is one of the most influential papers documenting the extent to which the effectiveness of aid depends on the quality of policies in the country.⁶ This paper gives in my view a reasonable, and probably representative, estimate of the growth effects of aid. The empirical specification is based on a panel dataset of four-year average growth rates spanning a large sample of developing countries. I decompose the variance of the dependent variable, growth, into portions explained by different groups of explanatory variables in the regression. Figure 5 summarizes the

⁶ While the results in this paper have been criticized for their lack of econometric robustness, it is important to note that the argument that aid works better in good policy and institutional environments is built on much more than this one paper. Burnside and Dollar (2004) provide updated cross-country empirical evidence in support of this interaction. Several papers have also documented that aid projects tend to work better in countries with good institutions (Isham and Kaufmann (1999), Dollar and Levin (2005)). And similar conclusions about the greater effectiveness of aid in good policy environments can be found in the case studies of aid and reform in Africa documented in Devarajan, Dollar, and Holmgren (2001).

results of this decomposition. The effect of aid on growth is captured by two variables, measuring the direct effect of aid, as well as the interaction of aid with their index of policy. While these terms are statistically significant, it is important to note that they account for just four percent of the variation in the dependent variable. By contrast, the direct effect of policies, as well as institutional quality, together account for some 17 percent of the variation in growth, and exogenous factors such as regional dummies and ethnic fractionalization, account for another 18 percent. In total, this still leaves 61 percent of the variation in growth unaccounted for.

My objective here is not to criticize this regression in particular, nor is it to criticize cross-country growth empirics in general. As noted above the Burnside-Dollar (2000) paper is rightly one of the most influential papers on the effects of aid on growth. Rather, this simple variance decomposition underscores the observation that growth depends on much more than just aid. Thus while aid can contribute to poverty alleviation, over the medium to long term where most of changes in poverty depend on growth, factors other than aid will be important as well for determining the pace of poverty reduction.

Finally, it should also be noted that aid can also have direct effects on poverty independent of any effect on aggregate growth. In terms of direct effects on income-poverty, such a redistributive channel is likely to be small – after all we have seen in the previous section that the contribution of relative income changes to poverty reduction has historically been small on average. But at the same time it is worth noting that a significant fraction of aid is – often very successfully – directed at attacking non-income forms of poverty in ways that may not have any significant impact on growth in the medium-run. For example, a recent book from the Center for Global Development documents 17 very successful public health interventions in developing countries, often significantly financed by foreign aid.⁷ These include campaigns against riverblindness in Africa, tuberculosis in China, measles in Latin America, and many more. There are also many successful examples of aid financing improvements in education, particularly among those traditionally having less access to schooling. And a renewed emphasis on impact evaluation among aid donors is contributing to making these kind of interventions increasingly effective. These type of human capital improvements vividly illustrate how

⁷ Center for Global Development (2004).

aid can directly impact non-income forms of poverty without necessarily affecting growth over the medium run, or possibly even the long run.

4. Governance, Poverty Traps, and Growth

We have seen that growth is central to poverty reduction, and that while there is evidence that aid can raise growth, the share of the variation in growth accounted for by aid is relatively modest. This emphasizes the importance of non-aid determinants of growth for poverty reduction. It also points to the question of how well aid can help to support these other determinants of growth. As noted earlier, this opens the door to a very large set of issues that could possibly be discussed. Since I cannot be comprehensive, I would like to instead discuss two broad issues that I have been working on in my recent research: governance, and poverty traps.

Governance and Growth in Africa

A large literature over the past 10 years or so has provided compelling evidence that various dimensions of good governance – including the protection of property rights, the absence of corruption, and the existence of competent and effective governments -- have important causal effects on economic development. There is also evidence in support of the eminently plausible proposition that aid is more effective in countries with good governance. Against this background it is sobering to note that available data suggests that the quality of governance in Africa remains on average quite low, and that there is also no strong pattern of improvements in governance that would improve the situation.

Figure 6 gives a snapshot of African countries' standing on one cross-country measure of property rights protection or "rule of law" in 2004, taken from Kaufmann, Kraay and Mastruzzi (2004). The governance measure is on the vertical axis, and countries are organized by per capita income on the horizontal axis. Note that the per capita income variable has been rescaled to have mean zero and standard deviation of one, as does the governance indicator. Not surprisingly countries in Africa are concentrated on the extreme left side of the graph, corresponding to very low income levels. And all but a handful of the countries in the region fall below the median score of

zero on the governance indicator. This graph shows clearly that countries in Africa in general are poor and have weak governance.

However, in a recent paper, Jeffrey Sachs has argued that weak governance is not a major factor in Africa's poor growth performance (Sachs et. al. (2004)). The argument is that, once we control for per capita income, countries in sub-Saharan Africa do not have particularly poor governance indicators. A simple way to capture this argument is to put a line of best fit through the points in Figure 6 and ask where countries in Sub-Saharan Africa fall relative to this line. A somewhat surprising observation from this graph is that over half (27 out of 46) of the countries in the region actually fall above the line of best fit, shown in black. At first glance this seems to suggest that if we take levels of development into account, governance in the region is roughly what might be expected given low incomes. This line of argument rapidly leads to two strong conclusions. First, it suggests that weak governance in the region may not be a major obstacle to aid effectiveness. And second, by implicitly arguing that Africa's weak governance is the causal outcome of low income levels, it suggests that increases in income will automatically lead to improvements in governance and that direct interventions to improve governance may not be a great priority.

I do not think that either of these conclusions are appropriate, because they flow from an overly simplistic interpretation of the evidence in Figure 6. This interpretation of the graph is valid only to the extent that the line of best fit captures a causal relationship from higher income to better governance. But there is a large body of research which indicates that there is substantial causation in the other direction as well – better governance leads to higher incomes. This means that the simple correlation represented by the black line in Figure 6 will exaggerate the positive effects of income on governance because it also reflects the strong effect in the opposite direction, from governance to incomes. In order to compare governance in Sub-Saharan Africa to what might be expected given income levels, we therefore need to first isolate these two directions of causation.

The red and green lines in Figure 6 show two alternative estimates of the causal effect of income on governance. The upward-sloping one comes from Rigobon and Rodrik (2004). They study the causal relationships between per capita income,

democracy, rule of law, openness to international trade, and geography, using identification through heteroskedasticity to isolate the causal effects.⁸ As expected, the red line is substantially flatter than the ordinary least squares regression line, consistent with the intuition that the latter relationship overstated the true causal effect of incomes on governance. This flattening has important consequences for our conclusions about the quality of governance in Africa controlling for income levels. Once we isolate this much weaker effect of income on governance, we find that only 7 out of 46 countries in the region fall above the regression line: Ghana, Lesotho, Cape Verde, Namibia, .South Africa, Botswana, and Mauritius. In contrast, the vast majority of countries in Africa have governance that is worse than their income levels would predict.

The downward-sloping green line presents another estimate of the effect of income on governance, coming from Kaufmann and Kraay (2002). They use a different approach to identification and find a zero or even negative impact of income on governance. While this finding is somewhat extreme, it leads to the same conclusions regarding the quality of governance in Africa – now only 6 out of 46 countries in the region fall above the regression line, indicating governance levels better than what per capita incomes would predict. Overall this evidence suggests that we cannot conclude that governance is not a problem in Sub-Saharan Africa, even after taking into account the region’s low per capita income levels. If anything, the evidence suggests that, even after controlling for incomes, governance in the region is worse than one might expect.

In order to isolate the causal effects of income on governance in Figure 6, we also had to identify the opposite direction of causation, from governance to per capita incomes. Figure 7 shows estimates of the causal relationship in this direction, with the two alternative estimates based on the same two papers discussed above. Both papers find a statistically significant causal impact of governance on per capita incomes, although the magnitudes of the estimated effects differ substantially. From this graph it

⁸ We use their specification excluding democracy, which implies that a one standard deviation increase in log per capita GDP improves rule of law by 0.14 standard deviations. They use a different measure of rule of law for the mid-1990s taken from Knack and Keefer (1995). However, its correlation with our rule of law indicator is above 0.8, so we can reasonably use the estimated coefficient from this paper with our governance indicator, suitably standardized. Note also that in the system of equations estimated by Rigobon and Rodrik (2004) the conditional expectation of governance given per capita income also reflects the indirect effects of income on openness, which in turn affects the rule of law. However, these estimated indirect effects are so small that our conclusions are essentially unaffected by ignoring them.

should be clear that while governance matters, it surely is not the only thing that matters as there is substantial dispersion around the fitted relationships. And for Africa in particular, it is again striking that the majority of countries in the region fall below the regression lines capturing the effects of governance on income. For the Rigobon and Rodrik (2004) estimate, 40 out of 46 countries fall below the line, and for the Kaufmann and Kraay (2002) estimate, 30 out of 46 countries fall below the line. This observation emphasizes the fact that low income levels in Africa are not solely attributable to on average weak governance performance, and that other factors such as Africa's difficult geography, its dependence on natural resources, the prevalence of civil conflict, and many other factors play a role. In short, I do not want to conclude from these graphs that governance is the *only* thing that matters. Rather I want to emphasize that governance does matter for development, and the causal effect of incomes on governance is sufficiently modest that it does not seem appropriate to "discount" the governance performance of the region simply because it is poor.

Finally it is useful to briefly examine available evidence on trends in governance in Africa. It is worth noting at the outset that by some measures, notably democratic accountability, there have been substantial improvements over the past 25 years. Figure 8 shows that there have been sharp increases in the proportion of countries with competitive elections, and this increase has been more dramatic in Africa than elsewhere, albeit from a lower base.⁹ Currently just over half of countries in Africa have chief executives installed as a consequence of competitive elections, a proportion similar to that in the rest of the developing world.

There is however less evidence of clearcut improvements of governance in the region over the more recent period for which we have more data. **Figure 9** reports a comparison of six different dimensions of governance between 1996 and 2004. The governance indicators are part of an ongoing project at the World Bank to measure governance, and capture perceptions of "voice and accountability", "political instability", "government effectiveness", "regulatory quality", "rule of law", and "control of corruption". The governance measures are composite indicators that combine data from a large

⁹ In this graph elections are deemed to be "competitive" if the winner obtained less than 75 percent of the vote. The data come from Beck and others (2001).

number of sources of data on perceptions of the quality of governance.¹⁰ The graph reports the average (across all countries in the region) of the percentile rank of countries on each of the six indicators, in the two periods. The top bar for each indicator corresponds to 2004, while the bottom bar corresponds to 1996. All of the bars are just above 0.25, indicating that the typical country in Africa has a governance score just above the bottom quartile for the world as a whole. It is clear from this graph that there is no obvious trend towards better governance in the region. In fact, the averages show very slight declines in all six dimensions, although one should be careful not to conclude too much from these changes because of the non-trivial margins of error associated with the governance estimates.

Figure 10 goes below these regional averages to show trends in one dimension of governance, the rule of law or property rights protection that we saw earlier. The 2004 score is on the vertical axis, while the 1996 score is on the horizontal axis, so countries falling below the 45-degree line correspond to countries where perceptions of this dimension of governance have deteriorated since 1996.¹¹ Countries in Africa are again labelled, and the vertical line shows the margin of error (technically a 90 percent confidence interval) associated with the governance estimate for each country. For many countries this vertical line crosses the 45-degree line, suggesting that the change in governance since 1996 is small relative to the unavoidable margins of error that arise when measuring governance. But there are also countries that have registered non-trivial improvements (such as Mozambique, Madagascar, and Mali), while there are quite a few others where there have been non-trivial declines (such as Zimbabwe, Cote d'Ivoire, and Somalia). The point of this graph is to show that even over relatively short periods such as the eight years shown here, there are movements in governance in both directions underlying the fairly stable regional averages shown in **Figure 9**.

I have argued above that good governance has important direct impacts on growth, and also is an important determinant of the effectiveness of foreign aid. In light of Africa's on average poor performance in various dimensions of governance, and in

¹⁰ See Kaufmann, Kraay, and Mastruzzi (2004) for a description of the governance indicators for 1996-2002. The 2004 indicators shown here are currently being finalized and will be released soon.

¹¹ Note that the graph has been truncated at the high end in both dimensions in order to make it more legible. Countries in roughly the top quarter of governance in both periods are not shown as a result.

the absence of compelling evidence of an upward trend for the region as a whole, it seems important for the purposes of this conference to ask what role aid can play in improving governance or institutional quality. It is difficult to do justice to such an important question in the limited time available. And it is all the more difficult given how little we as yet know from an academic and policymaking perspective about the dynamics of institutional change. It does seem useful to at least mention a few specific interventions that are promising. The World Bank Institute, at the request of several governments around the world, has carried out in-depth governance diagnostic surveys in countries that have been very useful in sparking a process of debate and discussion over the specifics of institutional strength and weakness in that country. Another example is the Extractive Industries Transparency Initiative proposed by the British government in 2002, developing a set of standards of conduct in the management of natural resource revenues. As argued by Collier (2004) this could be part of a larger set of standards or codes of conduct that could usefully be deployed to improve policymaking and accountability in developing countries, just as a different set of standards has been useful among rich countries (such as standards applying to EU member states by mutual agreement).

Poverty Traps, Aid, and Growth

Poverty traps have captured the imagination of academics and development practitioners for many years. It is not hard to see why – there are very many plausible self-reinforcing mechanisms whereby countries, or individuals, that start out poor might remain poor. If saving rates, or technology, or other positive forces for growth are low precisely because countries are poor, then countries may find themselves trapped at low levels of development. The poverty trap view of Africa's underdevelopment has also been made forcefully in Sachs et. al. (2004).

Despite both the popularity and plausibility of poverty traps, there is relatively little empirical work testing for poverty traps, and much of this tends not to be very supportive of the poverty trap hypothesis. Some of this work is at the very reduced-form level. A number of papers have documented that the distribution of per capita incomes across countries is gradually becoming bimodal over the past 50 years, with a group of

countries clustering around a quite low income level.¹² Another type of reduced-form evidence comes from looking at the dynamics of individual incomes. Many models of poverty traps suggest that individuals receiving large negative income shocks may take a very long time to recover, and if incomes fall below a certain threshold, they may never recover. However, Lokshin and Ravallion (2004) carefully examine household data from Hungary and Russia, and conclude that there is no evidence of the kind of “threshold effects” associated with models of poverty traps.

The difficulty with this kind of reduced-form evidence is that it provides no guidance as to what underlying mechanism is generating the poverty trap. Without this information it is difficult to formulate an appropriate policy response. Several recent studies have looked for evidence of particular mechanisms generating poverty traps. One such mechanism has to do with financial market imperfections. If the up-front cost of starting a small business is large, and poor individuals cannot borrow to finance this investment, then they will be unable to reap the benefits of self-employment. McKenzie and Woodruff (2004) use detailed data on microenterprises in Mexico and document that the costs of starting such a small business are surprisingly small, averaging just two weeks’ income of a typical low-wage Mexican worker. This casts doubt on the idea that fixed costs combined with financial frictions are responsible for poverty traps.

Another possible mechanism is that productivity is low at low levels of development. This may be because it is difficult to reach minimum efficient scales of production, or because complementary investments in public goods such as infrastructure are inadequate in poor countries. Once these thresholds are crossed it is possible that productivity increases sharply, allowing countries to reach much higher income levels. Kraay and Raddatz (2005) embed this mechanism in a standard growth model and show that for this mechanism to generate a poverty trap, it must be the case that productivity increases implausibly sharply with the level of development. In particular they show that if this mechanism is at work, we should expect to see increasing returns to scale that are substantially larger than is ever seen in the large empirical literature on estimating production functions. And somewhat more directly,

¹² See Azariadis and Stachurski (2004) for links between models of poverty traps and this kind of empirical evidence, and Quah (1993a, 1993b, 1996, and 1997) for the evidence, and Kremer, Stock and Onatski (2001) for a critique. Bloom, Canning, and Sevilla (2003) also provide closely-related cross-country evidence.

McKenzie and Woodruff (2004) find in their Mexican data that returns to investment are very high even for very small enterprises.

Poverty traps might also arise because saving rates are low in poor countries. If many households live at the margins of subsistence, they will be unable to save very much. Public saving might also be low at low income levels because governments of very poor countries have difficulty with tax collection. These low saving rates may translate into sufficiently low investment rates that countries are unable to accumulate very large stocks of productive assets per capita. And if saving rates only begin to increase at much higher levels of development, then countries that start out poor may be stuck in a poverty trap. Kraay and Raddatz (2004) take this hypothesis seriously but find little evidence in support of it. From an empirical perspective, they find no evidence that saving rates increase sufficiently quickly with the level of development to generate a poverty trap in a standard growth model with exogenous saving. They also calibrate a growth model with subsistence consumption and find that the impact on saving and growth is substantial only for countries that start out very close to subsistence levels. The significant dispersion in per capita incomes even within a poor region such as Sub-Saharan Africa therefore implies that the role of subsistence consumption can only explain low saving and growth in just a few of the very poorest countries in the region.

There are also potential poverty traps based on self-reinforcing dynamics in the area of governance. There is for example evidence that civil wars are both a consequence and a cause of low income, creating the possibility of a conflict trap (Collier et. al. 2003). There are also reasons to believe that high levels of corruption create self-perpetuating expectations of future corruption. The role of such mechanisms in generating stable poverty traps in growth models is not yet fully studied. But these mechanisms are arguably more plausible than some of the others discussed here.

What does all this imply for foreign aid? If saving or technological poverty traps were important, it would be likely that large-scale increases in aid would be necessary in order to get countries across the relevant thresholds and set them on sustained growth paths. But there are at least two reasons to be skeptical of such an argument. First, we have seen that the direct evidence for such traps is not very compelling. Second, most of the empirical evidence on the growth impacts of aid suggests that there are

diminishing, not increasing returns to aid. The implications for aid also depend on the mechanisms generating poverty traps in Africa. If for example civil conflict or corruption-related poverty traps are important, then large increases in financial assistance might actually be counterproductive, increasing incentives and opportunities for corruption and conflict. As argued in Collier (2004) tackling these underlying disfunctions directly must be done in parallel with any large increases in aid. Overall, however, the state of the empirical evidence should discourage us from any strong expectation that sufficiently large amounts of aid are likely to trigger sustainable growth booms as countries escape from poverty traps.

5. Conclusions

My task in this presentation was to discuss the links between aid and poverty reduction. As I have discussed in some detail, there is an abundance of evidence in support of the proposition that growth is central to poverty reduction. In light of this, the direct growth effects of aid create potentially a strong channel from aid to poverty reduction. But at the same time it should be remembered that cross-country differences in aid account for only a small share of the cross-country differences in growth performance, and for good reason. As Steve Radelet has discussed, not all types of aid should be expected to raise growth. And of course, growth depends on much more than just aid, so a better understanding of how aid can support these other drivers of growth seems important to the discussion of aid and poverty reduction.

I have tried to argue that tackling governance problems in Africa must be part of the growth and poverty reduction agenda, and while we as yet know less than we would like to about the process of improving governance, there are interventions where aid can play a role. At the same time, I have also cautioned against the argument that large-scale increases in aid are essential to breaking a poverty trap in Africa. We do not have sufficient empirical evidence to conclude that such a large scaling up of aid will have disproportionate effects on economic growth. This is not to say that Africa does not need, or cannot effectively use, more development assistance. Rather we should not be appropriately cautious in our expectations of what the long-term growth effects of such aid may be.

References

- Azariadis, Costas, and John Stachurski, (2004), "Poverty Traps", Forthcoming in the Handbook of Economic Growth, Aghion and Durlauf eds.
- Beck, Thorsten, George Clarke, Alberto Groff, Philip Keefer and Patrick Walsh (2001). "New Tools in Comparative Political Economy: The Database of Political Institutions". *World Bank Economic Review*. 15(1):165-176.
- Bloom, David E, Canning, David and Sevilla, Jaypee, (2003). "Geography and Poverty Traps," *Journal of Economic Growth*, vol. 8(4), pages 355-78, December.
- Burnside, Craig and David Dollar (2000). "Aid, Policies, and Growth." *American Economic Review* 90(4):847-68.
- Burnside, Craig and David Dollar (2004). "Aid, Policies, and Growth: Revisiting the Evidence." World Bank Policy Research Working Paper 3251.
- Clemens, Michael, Steven Radelet, and Rikhil Bhavnani (2004). "Counting Chickens When They Hatch: The Short-Term Effect of Aid on Growth". Center for Global Development Working Paper.
- Center for Global Development (2004). "Millions Saved: Proven Successes in Global Health". Washington.
- Collier, Paul and others (2003). "Breaking the Conflict Trap: Civil War and Development Policy". Oxford, Oxford University Press.
- Collier, Paul (2004). "African Growth, Why a Big Push?". Manuscript, Oxford University.
- Datt, Gaurav and Martin Ravallion (1992). "Growth and Redistribution Components of Changes in Poverty Measures: A Decomposition with Applications to Brazil and India in the 1980s". *Journal of Development Economics*. 38(2):275-295.
- Devarajan, Shantayanan, David Dollar and Torgny Holmgren, eds (2001). *Aid and Reform in Africa*. Washington, The World Bank.
- Dollar, David and Aart Kraay (2002). "Growth is Good for the Poor". *Journal of Economic Growth*. 7:195-225.
- Dollar, David and Victoria Levin (2005). "Sowing and Reaping". Manuscript, The World Bank.
- Easterly, William (1999). "Life During Growth". *Journal of Economic Growth*. 4(3):239-76.
- Isham, Jonathan and Daniel Kaufmann (1999). "The Forgotten Rationale for Policy Reform: The Performance of Investment Projects." *Quarterly Journal of Economics* 114(1):149-184.

- Kakwani, Nanak (2000). "What Is Pro-Poor Growth?". *Asian Development Review*. 18(1): 1-16.
- Kaufmann, Daniel and Aart Kraay (2002). "Growth Without Governance". *Economia*. 3(1):169-215.
- Kaufmann, Daniel, Aart Kraay and Massimo Mastruzzi (2004). "Governance Matters III: Governance Indicators for 1996, 1998, 2000, and 2002". *World Bank Economic Review*. 18:253-287.
- Kraay, Aart (2005). "When Is Growth Pro-Poor? Cross-Country Evidence". *Journal of Development Economics*, forthcoming.
- Kraay, Aart and Claudio Raddatz (2005). "Poverty Traps, Aid, and Growth". Manuscript, The World Bank. (Background paper for Global Monitoring Report 2005).
- Lokshin, Michael and Martin Ravallion (2004). "Household Income Dynamics in Two Transition Economies". *Studies in Nonlinear Dynamics and Econometrics*. 8(3) Article 4.
- McKenzie, David and Christopher Woodruff (2004). "Is There an Empirical Basis for Poverty Traps in Developing Countries?". Manuscript, Stanford and UCSD.
- Quah, Danny T, (1996). "Twin Peaks: Growth and Convergence in Models of Distribution Dynamics," *Economic Journal*, 106 (127), pages 1045-55.
- Quah, Danny T, (1993a), "Empirical Cross-Section Dynamics in Economic Growth". *European Economic Review* 37, pp. 426–434.
- Quah, Danny T, (1993b), "Galton's Fallacy and Tests of the Convergence Hypothesis". *The Scandinavian Journal of Economics* 95, pp. 427–443.
- Quah, Danny T, (1997), "Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs". *Journal of Economic Growth* 2 (1997), pp. 27–59.
- Ravallion, Martin and Shaohua Chen (2003). "Measuring Pro-Poor Growth". *Economics Letters*. 78:93-99.
- Rigobon, Roberto and Dani Rodrik (2004). "Rule of Law, Democracy, Openness, and Income: Estimating the Interrelationships. Manuscript. MIT and Kennedy School
- Sachs, Jeffrey, John W. McArthur, Guido Schmidt-Traub, Margaret Kruk, Chandrika Bahadur, Michael Faye, and Gordon McCord, (2004), "Ending Africa's Poverty Trap", *Brooking Papers on Economic Acitivity*, 2004 (1).

Table 1: Variance Decompositions of Changes in Poverty

	Share of Variance in Changes in Poverty Due to Growth Component	Share of Variance of Growth Component Due to Growth
<i>Long Spells (Average Length = 10 years)</i>		
Headcount	0.97	0.89
Poverty Gap	0.79	0.92
Squared Poverty Gap	0.68	0.92
Watts	0.72	0.92
<i>All Spells (Average Length = 3 years)</i>		
Headcount	0.70	0.91
Poverty Gap	0.53	0.92
Squared Poverty Gap	0.43	0.92
Watts	0.46	0.92

Figure 1: Growth and Poverty Reduction

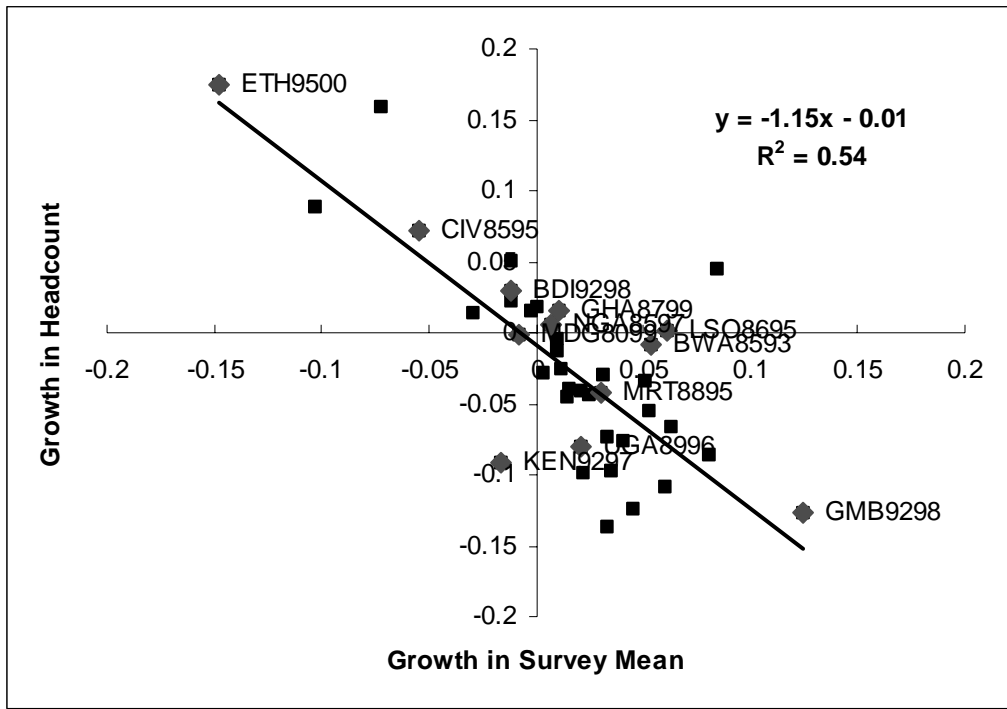


Figure 2: Decomposing Changes in Poverty: Growth vs Distribution




Figure 3: Decomposing Growth Component of Changes in Poverty


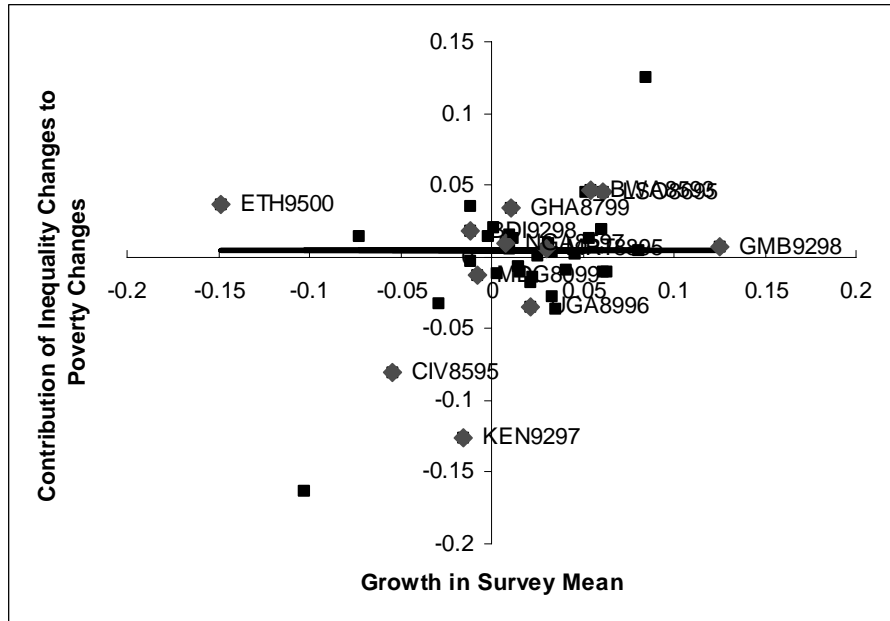


Figure 4: Growth and Poverty Reduction in Sub-Saharan Africa

Contribution of Inequality Changes



Contribution of Low Sensitivity of Poverty to Growth

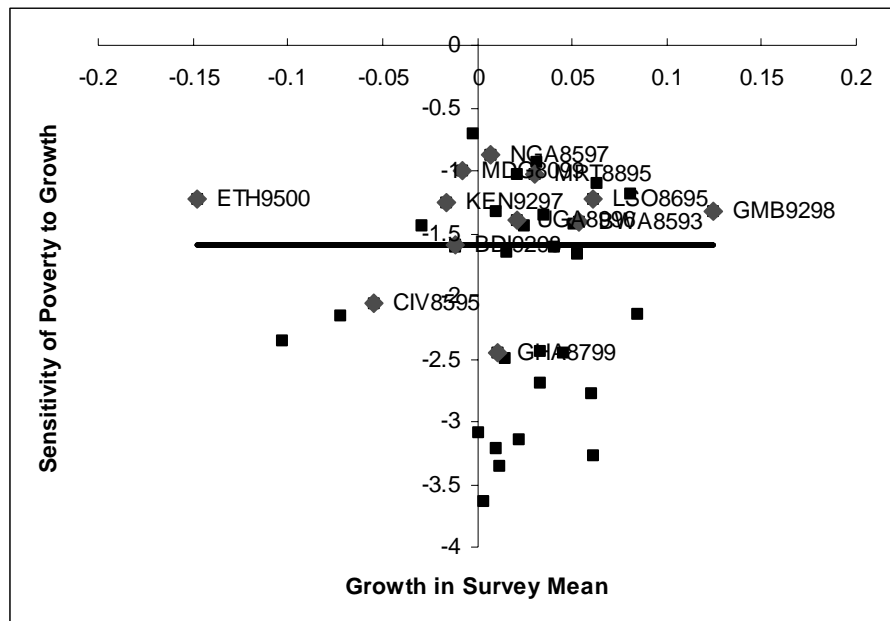


Figure 5: Estimated Contribution of Aid to Growth

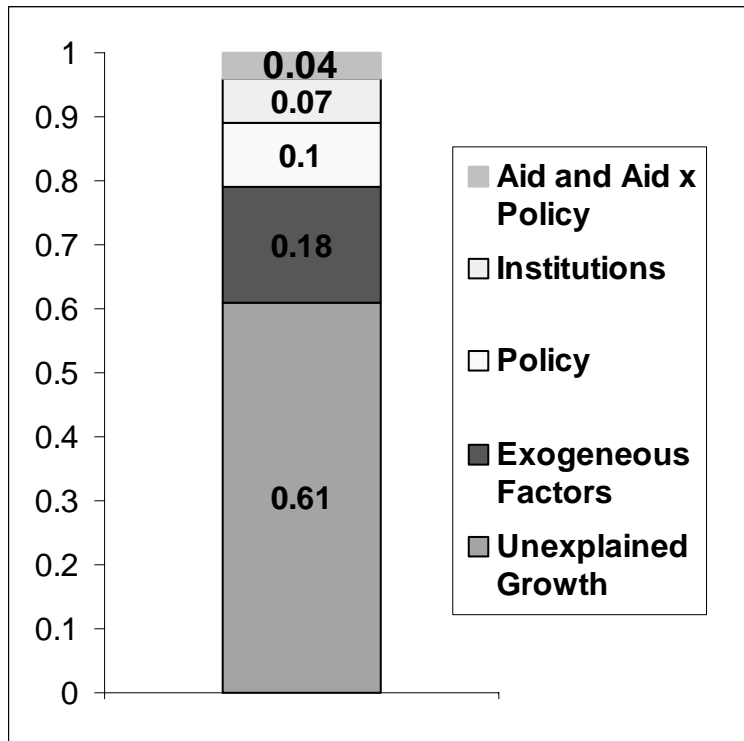


Figure 6: Governance and Growth in Sub-Saharan Africa





Figure 7: Effects of Governance on Income



**Figure 8: Improvements in Democratic Accountability Worldwide
(Proportion of Countries With Competitive Elections)**





Figure 9: Trends in Governance in Sub-Saharan Africa



Source: D. Kaufmann, A. Kraay and M. Mastruzzi, 2003: Governance Matters III: Governance Indicators for 1996-2002 (<http://www.worldbank.org/ubi/governance/pubs/govmatters3.html>)

Figure 10: Trends in Rule of Law 1996-2004

